arXiv:1701.07810v3 [cs.IR] 28 Feb 2017

Learning to Effectively Select Topics For Information Retrieval Test Collections

Mucahid Kutlu[1], Tamer Elsayed[1], and Matthew Lease[2]

[1]Dept. of Computer Science and Engineering, Qatar University, Qatar
[2]School of Information, University of Texas at Austin, USA

March 1, 2017

Abstract

While test collections are commonly employed to evaluate the effectiveness of information retrieval (IR) systems, constructing these collections has become increasingly expensive in recent years as collection sizes have grown ever-larger. To address this, we propose a new *learning-to-rank* topic selection method which reduces the number of search topics needed for reliable evaluation of IR systems. As part of this work, we also revisit the deep vs. shallow judging debate: whether it is better to collect many relevance judgments for a few topics or a few judgments for many topics. We consider a number of factors impacting this trade-off: how topics are selected, topic familiarity to judges, and how topic generation cost may impact both budget utilization and the resultant quality of judgments. Experiments on NIST TREC Robust 2003 and Robust 2004 test collections show not only our method's ability to reliably evaluate IR systems using fewer topics, but also that when topics are intelligently selected, deep judging is often more cost-effective than shallow judging in achieving the same level of evaluation reliability. Topic familiarity and construction costs are also seen to notably impact the evaluation cost vs. reliability tradeoff and provide further evidence supporting deep judging in practice.

1 Introduction

Test collections provide the cornerstone for system-based evaluation of information retrieval (IR) algorithms in the Cranfield paradigm [17]. Such test collections are vital to ensuring continuing field progress by enabling A/B testing of new search algorithms and community benchmarking. However, realistic evaluation requires testing IR systems at the scale of information that is to be searched in practice, and this has become increasingly problematic as today's collection sizes have grown ever-larger. In particular, because larger collections tend to contain more relevant documents (which retrieval systems might return in response to a given search query), human assessors must judge the relevance of ever-more documents for each search topic. Moreover, if insufficient documents are judged, evaluation findings could be compromised [10]. Consequently, a key challenge in IR today is to devise new evaluation methods to reduce the cost of constructing test collections while continuing to ensure the validity of findings.

One direction of research to lower costs has sought to reduce the number of topics for which human relevance judging is needed at all. NIST TREC test collections have traditionally used 50 search topics, with Buckley and Voorhees [11] reporting that at least 25 topics are needed for stable evaluation, with 50 being better. However, Zobel [60] showed that one set of 25 topics predicted relative performance of systems fairly well on a different set of 25 topics. Guiver, Mizzaro, and Robertson [22] conducted a systematic study showing that evaluating IR systems using the "right" subset of topics yields very similar results vs. evaluating systems over all topics. However, they did not propose a method to find such an effective topic subset in practice. NIST employs a simple and effective but costly topic selection process which includes initial judgments for each topic and manual selection from candidate topics [54]. Hosseini et al. [27] recently proposed an iterative method to automatically find effective topic subsets to lower evaluation costs while preserving evaluation reliability.

For a given evaluation budget, another related question is whether it is better to collect many relevance judgments for a few topics – i.e., *Narrow and Deep* (NaD) judging – or a few relevance judgments for many topics – i.e., *Wide and Shallow* (WaS) judging. Intuitively, since people search for a wide variety of topics with a wide variety of queries, it seems that realistic IR evaluation ought to evaluate systems across a similarly wide variety of search topics and queries. Empirically, large variance in search accuracy is often observed for the same system across different topics [7], suggesting merit in sampling many diverse topics for evaluation in order to achieve stable evaluation among systems being compared. Studies on this issue have led to a fairly consistent finding that WaS judging tends to provide more stable evaluation for the same human effort vs. NaD judging [48, 16, 9]. However, this finding does not hold in all cases. For example, Carterette et al. [14] achieve the same reliability using 250 topics with 20 judgments per topic (5000 judgments in total) as 600 topics with 10 judgments per topic (6000 judgments in total). Moreover, if we account for the time required to create each topic, the relative expense of using more topics is further increased.

In this work, we investigate how to automatically select informative search topics for IR evaluation [22, 27]. We begin with our first research question:

- **RQ-1: Intelligent Topic Selection.** *How can we minimize the number of topics (hence the cost) required for reliable evaluation of IR systems?* Toward this end, we propose a novel *learning-to-rank* (L2R) approach in which topics are selected iteratively via a greedy method which optimizes accurate ranking of systems (Section 4.3). Our L2R model utilizes 63 features representing the interaction between topics and ranking of systems (Section 4.3.1). In order to train our model, we describe a method to automatically generate useful training data from existing test collections (Section 4.3.2). We evaluate our approach on NIST TREC Robust 2003 [53] and Robust 2004 [52] test collections, comparing our approach to recent prior work [27] and *random* topic selection (Section 5). Results show consistent improvement over baselines, with greater relative improvement as fewer topics are used.

As mentioned above, we also revisit the deep vs. shallow judging debate, investigating how intelligent selection of topics impacts optimization of evaluation budget. In relation to past studies, our work is distinguished by investigating a wider range of factors impacting this tradeoff than previously considered:

- **RQ-2: NaD vs. WaS.** Prior studies of NaD vs. WaS judging have assumed random selection of topics. *How does intelligent selection of topics (by itself) impact evaluation budget optimization in NaD vs. WaS judging?* We find that NaD judging often achieves greater evaluation reliability for the same budget than WaS judging when topics are selected intelligently.

- **RQ-3: Judging Speed.** Past comparisons between NaD vs. WaS judging have typically assumed constant judging speed [16, 48]. However, data reported by Carterette et al. [15] suggests that assessors may judge

documents faster as they judge more documents for the same topic (likely due to increased topic familiarity). *Assuming WaS judging leads to slower judging speed than NaD judging, how does this impact optimization of evaluation budget?* Because we can collect more judgments in the same amount of time with NaD vs. WaS judging, NaD judging achieves greater relative evaluation reliability than shown in prior studies which did not consider this judging speed benefit of deep judging.

- **RQ-4: Topic generation Time.** Prior NaD vs. WaS studies have typically ignored non-judging costs involved in test collection construction. While Carterette et al. [14] consider topic generation time, the 5-minute time they assumed is roughly two orders of magnitude less than the 4 hours NIST has traditionally taken to construct each topic [1]. *How do historically large topic generation costs impact optimization of evaluation budget in the NaD vs. WaS tradeoff?* We find that WaS judging is preferable than NaD judging for short topic generation times (specifically $\leq$5 minutes in our experiments). However, as the topic generation cost further increases, NaD judging becomes increasingly preferable over WaS judging in optimizing evaluation budget.

- **RQ-5: Judging Error.** Several studies have reported calibration effects impacting the decisions and consistency of relevance assessors [47, 49]. While NIST has traditionally included an initial "burn-in" judging period as part of topic generation and formulation [1], we posit that drastically reducing topic generation time (e.g., from 4 hours [1] to 2 minutes [15]) could negatively impact topic quality, leading to less well-defined topics and/or calibrated judges, and thereby less reliable judgments. As suggestive evidence, McDonnel et al. [33] report high judging agreement in reproducing a "standard" NIST track, but high and inexplicable judging disagreement on TREC's Million Query track [15], which lacked any burn-in period for judges and had far shorter topic generation times. *Assuming short topic generation times reduce consistency of relevance judgments, how does resultant judging error impact evaluation budget optimization in balancing deep vs. shallow judging?* To investigate this, we simulate increased judging error as a function of lower topic generation times. We find that it is better to invest a portion of our evaluation budget to increase quality of topics, instead of collecting more judgments for low-quality topics. This also makes NaD judging preferable over WaS judging in many cases, due to increased topic generation cost.

The remainder of this article is organized as follow. Section 2 reviews the related work on topic selection and topic set design. Section 3 formally defines the topic selection problem. In Section 4, we explain our proposed L2R-based approach in detail. Section 5 presents our experimental evaluation. Finally, Section 6 summarizes the contributions of our work and suggests potential future directions.

2 Related Work

Constructing test collections is expensive in human effort required. Therefore, researchers have proposed a variety of methods to reduce the cost of creating test collections. Proposed methods include: developing new evaluation measures and statistical methods for the case of incomplete judgments [5, 12, 41, 59, 58], finding the best sample of documents to be judged for each topic [18, 13, 29, 36, 39], inferring relevance judgments [6], topic selection [27, 28, 35, 22], evaluation with no human judgments [38, 50], crowdsourcing [4, 21], and others. The reader is referred to [37] and [46] for a more detailed review of prior work on methods for low-cost IR evaluation.

2.1 Topic Selection

To the best of our knowledge, Mizzaro and Robertson [35]'s study was the first seeking to select the best subset of topics for evaluation. They first build a system-topic graph representing the relationship between topics and IR systems, then run the HITS algorithm on it. They hypothesize that topics with higher hubness scores would be better to distinguish between systems. However, Robertson [40] experimentally shows that their hypothesis is not true.

Guiver, Mizzaro, and Robertson [22] experimentally show that if we choose the right subset of topics, we can achieve a ranking of systems that is very similar to the ranking when we employ all topics. However they do not provide a solution to find the right subset of topics. This study has motivated other researchers to investigate this problem: Berto, Mizzaro, and Robertson [8] stress generality and shows that a carefully selected good subset of topics to evaluate a set of systems can be also adequate to evaluate a different set of systems. Hauff et al. [23] report that using the easiest topics based on Jensen-Shannon Divergence approach does not work well to reduce the number of topics. Hosseini et al. [28] focus on selecting the subset of topics to extend an existing collection in order to increase its re-usability. Kazai and Sung [30] reduce the cost for preference based IR evaluation.

The closest study to our own is [27], which employs an adaptive algorithm for topic selection. It selects the first topic randomly. Once a topic is selected, the relevance judgments are acquired and used to assist with the selection of subsequent topics. Specifically, in the following iterations, the topic that is predicted to maximize the current Pearson correlation is selected. In order to do that, they predict relevance probabilities of *qrels* for the remaining topics using a Support Vector Machine (SVM) model trained on the judgments from the topics selected thus far. Training data is extended at each iteration by adding the relevance judgments from each topic as it is selected in order to better select the next topic.

Further studies investigate topic selection for other purposes, such as creating low-cost datasets for training learning-to-rank algorithms [34], system rank estimation [24], and selecting training data to improve supervised data fusion algorithms [32]. These studies do not consider topic selection for low-cost evaluation of IR systems.

2.2 How Many Topics Are Needed?

Past work has investigated the ideal size of test collections and how many topics are needed for a reliable evaluation. While traditional TREC test collections employ 50 topics, a number of researchers claim 50 topics are not sufficient for a reliable evaluation [29, 55, 51]. Many researchers report that wide and shallow judging is preferable than narrow and deep judging [48, 16, 9], without providing the balance between pool depth and the number of topics. Carterette et al. [14] experimentally compare deep vs. shallow judging in terms of budget utilization. They find that 20 judgments with 250 topics is the most cost-effective in their experiments. Urbano, Marrero, and Mart n [51] measure the reliability of TREC test collections with regard to generalization and conclude that the number of topics needed for reliable evaluation varies across different tasks.

In order to calculate the number of topics required, Webber, Moffat, and Zobel [57] propose adding topics iteratively until desired statistical power is reached, while Sakai proposes methods based on two-way ANOVA [45], confidence interval [43], t test and one-way ANOVA [42]. In another work, Sakai also compares his previous methods and provides guidelines for test collection design with a given fixed budget [44]. While these works focus on calculating the number of topics required, our work focuses on how to select the best topic set for a given size in order to maximize the reliability of evaluation. We also introduce further considerations impacting the debate over shallow vs. deep judging: familiarization of users to topics, and the effect of topic generation costs on the budget utilization and the quality of judgments for each topic.

2.3 Topic Familiarity vs. Judging Speed

Carterette et al. [15] report that as the number of judgments per topic increases (when collecting 8, 16, 32, 64 or 128 judgments per topic), the median time to judge each document decreases respectively: 15, 13, 15, 11 and 9 seconds. This suggests that assessors become more familiar with the topic as they judge more documents for it, and this greater familiarity yields greater judging speed. However, prior work comparing deep vs. shallow judging did not consider this, instead assuming that judging speed is constant regardless of judging depth. Consequently, our experiments in Section 5.4 revisit this question, considering how faster judging with greater judging depth per topic may impact the tradeoff between deep vs. shallow judging in maximizing evaluation reliability for a given budget in human assessor time.

2.4 Topic Generation Cost vs. Judging Consistency

Past work has utilized a variety of different processes to generate search topics when constructing test collections. These different processes explicitly or implicitly enact potentially important tradeoffs between human effort (i.e. cost) vs. quality of the resultant topics generated by each process. For example, NIST has employed a relatively costly process in order to ensure creation of very high quality topics [1]:

> For the traditional ad hoc tasks, assessors generally came to NIST with some rough ideas for topics having been told the target document collection. For each idea, they would create a query and judge about 100 documents (unless at least 20 of the first 25 were relevant, in which case they would stop at 25 and discard the idea). From the set of candidate topics across all assessors, NIST would select the final test set of 50 based on load-balancing across assessors, number of relevant found, eliminating duplication of subject matter or topic types, etc. The judging was an intrinsic part of the topic development routine because we needed to know that the topic had sufficiently many (but not too many) relevant in the target document set. (These judgments made during the topic development phase were then discarded. Qrels were created based only on the judgments made during the official judgment phase on pooled participant results.) We used a heuristic that expected one out of three original ideas would eventually make it as a test set topic. Creating a set of 50 topics for a newswire ad hoc collection was budgeted at about 175-225 assessor hours, which works out to about 4 hours per final topic.

In contrast, the TREC Million Query (MQ) Track used a rather different procedure to generate topics. In the 2007 MQ Track [3], 10000 queries were sampled from a large search engine query log. The assessment system showed 10 randomly selected queries to each assessor, who then selected one and converted it into a standard TREC topic by back-fitting a topic description and narrative to the selected query. Carterette et al. [14] report that median time of generating a topic was roughly 5 minutes. In the 2008 MQ Track [2], assessors could refresh list of candidate 10 queries if they did not want to judge any of the candidates listed. Carterette et al. [15] report that *median* time for viewing a list of queries is 22 seconds and back-fitting a topic description is 76 seconds. On average, each assessor viewed 2.4 lists to generate each topic. Therefore, the cost of generating a topic is roughly $2.4 * 22 + 76 \approx 129$ seconds, or 2.1 minutes.

The examples above show a vast range of topic creation times: from 4 hours to 2 minutes per topic. Therefore, in Section 5.4, we investigate deep vs. shallow judging when cost of generating topics is also considered.

In addition to considering topic construction time, we might also consider whether aggressive reduction in topic creation time might also have other unintended, negative impacts on topic quality. For example, Scholer et al. [49]

report calibration effects change judging decisions as assessors familiarize themselves with a topic. Presumably NIST's 4 hour topic creation process provided judges ample time to familiarize themselves with a topic, and as noted above, judgments made during the topic development phase were then discarded. In contrast, it seems MQ track assessors began judging almost immediately after selecting a query for which to back-fit a topic, and with no initial topic formation period for establishing the topic and discarding initial topics made during this time. Further empirical evidence suggesting quality concerns with MQ track judgments was also recently reported by McDonnell et al. [33], who describe a detailed judging process they employed to reproduce NIST judgments. While the authors report high agreement between their own judging and crowd judging vs. NIST on the 2009 Web Track, for NIST judgments from the 2009 MQ track, the authors and crowd judges were both consistent while disagreeing often with NIST judges. The authors also reported that even after detailed analysis of the cases of disagreement, they could not find a rationale for the observed MQ track judgments. Taken in sum, these findings suggest that aggressively reducing topic creation time may negatively impact the quality of judgments collected for that topic. For example, while an assessor is still formulating and clarify a topic for himself/herself, any judgments made at this early stage of topic evolution may not be self-consistent with judgments made once the topic is further crystallized. Consequently, in Section 5.4 we revisit the question of deep judging of few topics vs. shallow judging of many topics, assuming that low topic creation times may be also mean less consistent judging.

3 Problem Definition

In this section, we define our topic selection problem. We assume that we have a TREC-like setup: a data collection has already been acquired, a large pool of topics and ranked lists of IR systems for each topic are also available. Our goal is to select a certain number of topics from the topic pool such that evaluation with those selected topics yields the most similar ranking of the IR systems to the "*ground-truth*". We assume that the ground-truth ranking of the IR systems is the one when we use all topics in the pool for evaluation.

We can formulate this problem as follows. Let $T = \{t_1, t_2, ..., t_N\}$ denote the pool of N topics, $S = \{s_1, s_2, ..., s_K\}$ denote the set of K IR systems to be evaluated, and $R_{<S,T,e>}$ denote the ranking of systems in S when they are evaluated based on evaluation measure e over the topic set T (notation used in equations and algorithms is shown in **Table 1**). We aim to select a subset $P \subset T$ of M topics that *maximizes* the correlation (as a measure of similarity between two ranked lists) between the ranking of systems over P (i.e., considering only M topics and their corresponding relevance judgments) and the ground-truth ranking of systems (over T). Mathematical definition of our goal is as follows:

$$\max_{P \subset T,\ |P|=M} corr(R_{<S,P,e>}, R_{<S,T,e>}) \tag{1}$$

where *corr* is a ranking similarity function (*e.g.,* Kendall-τ [31]).

4 Proposed Approach

The problem we are tackling is challenging since we do not know the actual performance of systems (i.e., their performance when all topics are employed for evaluation) and we would like to find a subset of topics that achieves similar ranking to the *unknown* ground-truth.

Table 1: Notation used in equations and algorithms

Symbol	Name
T	Topic pool
S	IR systems participated to the pool of the corresponding test collection
$R_{<S,T,e>}$	Ranking of systems in S when they are evaluated based on evaluation measure e over the topic set T
N	Size of topic pool
M	Number of topics to be selected
D_{t_c}	The document pool for topic t_c
$L_{s_j(t_c)}$	The ranked list resulting from system s_j for the topic t_c

To demonstrate the complexity of the problem, let us assume that we obtain the judgments for all topic-document pairs (i.e., we know the ground-truth ranking). In this case, we have to check $\binom{N}{M}$ possibilities of subsets in order to find the optimal one (i.e., the one that produces a ranking that has the maximum correlation with the ground-truth ranking). For example, if $N = 100$ and $M = 50$, we need to check around 10^{29} subsets of topics. Since this is computationally intractable, we need an approximation algorithm to solve this problem. Therefore, we first introduce a greedy approach to select the best subset of topics when we have the judgments for all query-document pairs (Section 4.1). Subsequently, we discuss how we can employ this greedy approach in the lack of the relevance judgments (Section 4.2) and then introduce our L2R-based topic selection approach (Section 4.3).

4.1 Greedy Approach

We first explore an oracle greedy approach that selects topics in an iterative way when relevance judgments are already obtained. Instead of examining all possibilities, at each iteration, we select the '*best*' topic (among the currently non-selected ones) that, when added to the currently-selected subset of topics, will produce the ranking that has the maximum correlation with the ground-truth ranking of systems.

Algorithm 1 illustrates this oracle greedy approach. First, we initialize set of selected topics (P) and set of candidate topics to be selected ($\bar{P}$) (Line 1). For each candidate topic t in $\bar{P}$, we rank the systems over the selected topics P in addition to t ($R_{<S,P\cup\{t\},e>}$), and calculate the Kendall's τ achieved with this ranking (Lines 3-4). We then pick the topic achieving the highest Kendall-τ score among other candidates (Line 5) and update P and $\bar{P}$ accordingly (Lines 6-7). We repeat this process until we reach the targeted subset size M (Lines 2-7).

Algorithm 1 An Oracle Greedy Approach for Topic Selection

1: $P \leftarrow \emptyset$; $\bar{P} \leftarrow T$ ▷ selected set P is initially empty
2: **while** $|P| < M$ **do**
3: **for each** topic $t \in \bar{P}$ **do**
4: $\tau_t \leftarrow corr(R_{<S,P\cup\{t\},e>},\ R_{<S,T,e>})$
5: topic $t^* \leftarrow \underset{t\in\bar{P}}{max}(\tau_t)$ ▷ choose t^* yielding the best correlation
6: $P \leftarrow P \cup \{t^*\}$ ▷ add t^* to the selected set P
7: $\bar{P} \leftarrow \bar{P} - \{t^*\}$

While this approach has $O(M \times N)$ complexity (which is clearly much more efficient compared to selecting the optimal subset), it is also impractical due to leveraging the real judgments (which we typically do not have in advance) in order to calculate the ground-truth ranking and thereby Kendall-τ scores.

4.2 Performance Prediction Approach

One possible way to avoid the need for the actual relevance judgments is to predict the performance of IR systems using automatic evaluation methods (*e.g.,* [50, 38]) and then rank the systems based on their predicted performance. However, relying on predicted performance to represent the ground-truth is problematic simply because inaccurate predictions result in noisy *predicted* ground-truth, which causes the selection of topics that actually do not produce maximum correlation with the *real* ground-truth ranking.

One noteworthy study using this approach is Hosseini et al. [27]'s work which predicts relevance probability of document-topic pairs by employing an SVM classifier and selects topics in a greedy way similar to the approach in Algorithm 1. We use their selection approach as a baseline in our experiments presented in Section 5.

4.3 Proposed Learning-to-Rank Approach

In this work, we formulate the topic selection problem as a learning-to-rank (L2R) problem. In a typical L2R problem, we are given a query q and a set of documents D, and a model is learned to rank those documents in terms of relevance with respect to q. The model is trained using a set of queries and their corresponding labeled documents. In our context, we are given a set of currently-selected topic set P (analogous to the query q) and the set of candidate topics $\bar{P}$ to be selected from (analogous to the documents D), and we aim to train a model to rank the topics in $\bar{P}$ based on the expected effectiveness of adding each to P. The training samples used to train the model are tuples of the form $(P, t, corr(R_{<S,P\cup\{t\},e>}, R_{<S,T,e>}))$, where the measured correlation is considered the label of the topic t with respect to P. Notice that, in the training data, the correlation is computed using the true relevance judgments. This enables us to use the wealth of existing test collections to acquire data for training our model as explained in Section 4.3.2.

We apply this L2R problem formulation to the topic selection problem using our greedy approach. Instead of picking the topic that leads to the best correlation with the ground-truth among all candidate topics, we use the trained L2R model to rank the candidate topics and then pick the one at the first rank. The algorithm is illustrated in **Algorithm 2**. At each iteration, a feature vector v_t is computed for each candidate topic t in $\bar{P}$ using a feature extraction function f (Lines 3-4), detailed in Section 4.3.1. The candidate topics are then ranked using our learned

model (Line 5) and the topic in the first rank is picked (Line 6). Finally, the topic sets P and $\bar{P}$ are updated (Lines 7-8) and a new iteration is started, if necessary.

Algorithm 2 L2R-based Topic Selection

1: $P \leftarrow \emptyset$; $\bar{P} \leftarrow T$ ▷ selected set P is initially empty
2: **while** $|P| < M$ **do**
3: **for each** topic $t \in \bar{P}$ **do**
4: $v_t \leftarrow f(t, P, \bar{P})$ ▷ compute feature vector for t
5: $l \leftarrow L2R(\{(t, v_t)\} \mid \forall t \in \bar{P})$ ▷ apply L2R on feature vectors
6: topic $t^* \leftarrow first(l)$ ▷ choose t^* in the first rank
7: $P \leftarrow P \cup \{t^*\}$ ▷ add t^* to the selected set P
8: $\bar{P} \leftarrow \bar{P} - \{t^*\}$ ▷ Remove t^* from the candidate set $\bar{P}$

4.3.1 Features

In this section, we describe the features we extract in our L2R approach for each candidate topic. Hosseini et al. [27] mathematically show that, in the greedy approach, the selected topic at each iteration should be different from the already-selected ones (i.e., topics in P) while being a representative of the non-selected ones (i.e., topics in $\bar{P}$). Therefore, the extracted set of features should cover the candidate topic as well as the two sets P and $\bar{(P)}$. Our main goal is to capture the interaction between the topics and the IR systems in addition to the diversity between the IR systems in terms of their retrieval results.

We define two types of feature sets; the first (called *topic-based*) is extracted from an individual topic, and the second (called *set-based*) is extracted from a set of topics by *aggregating* the topic-based features extracted from each of those topics.

The topic-based features include 7 features that are extracted for a given candidate topic t_c and listed in **Table 2**. For a given set of topics (e.g., currently-selected topics P), we extract the set-based features by computing both average and standard deviation of each of the 7 topic-based features extracted from all topics in the set. This gives us 14 set-based features that can be extracted for a set of topics. We compute those 14 features for each of the following set of topics:

- currently-selected topics (P)
- not-yet-selected topics ($\bar{P}$)
- selected topics with the candidate topic ($P \cup \{t_c\}$)
- not-selected topics excluding the candidate topic ($\bar{P} - \{t_c\}$)

$\bar{P} - \{t_c\}$ and $P \cup \{t_c\}$ are used to represent, more clearly, how much the selected topics and not-selected topics will be affected in case we select the candidate topic.

Eventually, we have 63 features for each data record representing a candidate topic: $14 \times 4 = 56$ features for the above groups + 7 topic-based features. We describe the seven topic-based features that are the core of the entire feature set.

Table 2: Topic-based Features

Feature	Description
$f_{\bar{w}}$	Average sampling weight of documents
f_{σ_w}	Standard deviation of weight of documents
$f_{\bar{\tau}}$	Average τ score for ranked lists pairs
f_{σ_τ}	Standard deviation of τ scores for ranked lists pairs
$f_{\$}$	Judgment cost of the topic
$f_{\sigma_\$}$	Standard deviation of judgment costs of system pairs
f_{NSD}	Standard deviation of estimated performance of systems

- **Average sampling weight of documents ($f_{\bar{w}}$)**: In the statAP sampling method [39], a weight is computed for each document based on where it appears in the ranked lists of all IR systems. Simply, the documents at higher ranks get higher weights. The weights are then used in a non-uniform sampling strategy to sample more documents relevant to the corresponding topic. To compute this feature, we leverage that idea and compute the average sampling weight of all documents that appear in the pool of the candidate topic t_c as follows:

$$f_{\bar{w}}(t_c) = \frac{1}{|D_{t_c}|} \sum_{d \in D_{t_c}} w(d, S) \tag{2}$$

  where D_{t_c} is the document pool for topic t_c and $w(d, S)$ is the weight of document d over the IR systems S. High $f_{\bar{w}}$ values mean that the systems have common documents at higher ranks for the corresponding topic, whereas lower $f_{\bar{w}}$ values indicate that the systems return significantly different ranked lists or only the documents at lower ranks are in common.

- **Standard deviation of weight of documents (f_{σ_w})**: Similar to $f_{\bar{w}}$, we also compute the standard deviation of the sampling weights of documents for the candidate topic as follows.

$$f_{\sigma_w}(t_c) = \sigma\{w(d, S) \mid \forall d \in D_{t_c}\} \tag{3}$$

- **Average τ score for ranked lists pairs ($f_{\bar{\tau}}$)**: This feature computes Kendall's τ correlation between ranked lists of each pair of the IR systems and then take the average (as shown in Equation 4) in order to capture the diversity of the results of the IR systems. The depth of the ranked lists is set to 100. In order to calculate Kendall's τ score, the documents that appear in one list but not in the other are concatenated to the other list so that both ranked lists contain the same documents. If there are multiple documents to be concatenated, the order of the documents in the ranked list is preserved during concatenation. For instance, if system A returns documents {a,b,c,d} and system B returns {e,a,f,c} for a topic, then, the concatenated ranked lists of A and B are {a,b,c,d,e,f} and {e,a,f,c,b,d}, respectively.

$$f_{\bar{\tau}}(t_c) = \frac{1}{2|S| - 1} \sum_{i=1}^{|S|-1} \sum_{j=i+1}^{|S|} corr(L_{s_i(t_c)}, L_{s_j(t_c)}) \tag{4}$$

  where $L_{s_j(t_c)}$ represents the ranked list resulting from system s_j for the topic t_c.

- **Standard deviation of τ scores for ranked lists pairs (f_{σ_τ})**: This feature computes the standard deviation, in contrast to the average, of the τ scores of the pairs of ranked lists of systems as follows:

$$f_{\sigma_\tau}(t_c) = \sigma\{corr(L_{s_i(t_c)}, L_{s_j(t_c)}) \mid \forall i, j \leq |S|, i \neq j\} \tag{5}$$

- **Judgment cost of the topic ($f_\$$)**: This feature estimates the cost of judging the candidate topic as the number of documents in the pool at a certain depth. If IR systems return many different documents, then the judging cost increases; otherwise, it decreases due to having many documents in common. We set pool depth to 100 and normalize costs by dividing by the maximum possible cost (i.e., $100 \times |S|$). The feature is computed as follows:

$$f_\$(t_c) = \frac{D_{t_c}}{|S| \times 100} \tag{6}$$

- **Standard deviation of judgment costs of system pairs ($f_{\sigma_\$}$)**: The judgment cost depends on systems participated to the pool. In this feature, we construct the pool separately for each pair of systems and compute the standard deviation of the judgment cost across pools as follows:

$$f_{\sigma_\$}(t_c) = \sigma\{|L_{s_i(t_c)} \cup L_{s_j(t_c)}| \mid \forall i, j \leq |S|, i \neq j\} \tag{7}$$

- **Standard deviation of estimated performance of systems (f_{NSD})**: We finally compute standard deviation of the estimated estimated performances of the IR systems for the topic t_c using NSD query performance predictor [19] (See Equation 8). NSD predictor is typically used to estimate the performance of a single system and it is affected by the range of the retrieval scores of retrieved documents. Therefore, we normalize the document scores using min-max normalization before computing the predictor.

$$f_{NSD}(t_c) = \sigma\{NSD(s_i, t_c) \mid \forall i \leq |S|\} \tag{8}$$

  where $NSD(s_i, t_c)$ is the performance predictor applied on system s_i given topic t_c.

4.3.2 Generating Training Data

In our proposed L2R approach, we rank the topics based on their effectiveness when added to a certain set of topics. This makes obtaining the training data for the model a challenging task. First, there are countless number of possible scenarios (*i.e.*, different combinations of topic sets) that we can encounter during the topic selection process. Second, the training data should specify which topic is more preferable for a given scenario.

We developed a method to generate training data by leveraging existing test collections for which we have the relevance judgments and resulted ranked lists of documents for several IR systems (*e.g.*, TREC test collections). In our method, we first simulate a scenario where a number of topics have already been selected, then we rank the rest of the topics based on the correlation with the ground-truth ranking when each is added to the currently-selected topics. We repeat this simulation scenario multiple times and vary the number of already-selected topics to capture more diverse scenarios. The algorithm of generating training data from one test collection is given in **Algorithm 3**. Theoretically, the algorithm can be applied to several test collections to generate larger training data.

Algorithm 3 Generating Training Data from one test collection of T topics and S IR systems

1: $R^* \leftarrow R_{<S,T,e>}$ ▷ Determine the ground-truth ranking of systems
2: **for** i = 0 to N - 1 **do** ▷ Vary the size of already-selected topic set
3: **for** j = 1 to W **do** ▷ Repeat the process W times
4: $P \leftarrow random(T, i)$ ▷ Randomly select i topics
5: $\bar{P} \leftarrow T - P$
6: **for each** topic $t \in \bar{P}$ **do**
7: $R_t \leftarrow R_{<S,P \cup t,e>}$
8: $\tau_t \leftarrow corr(R_t, R^*)$
9: $\tau_{max} \leftarrow \underset{t \in \bar{P}}{max}(\tau_t)$; $\tau_{min} \leftarrow \underset{t \in \bar{P}}{min}(\tau_t)$
10: **for each** topic $t \in \bar{P}$ **do**
11: $v_t \leftarrow f(t, P, \bar{P})$ ▷ Compute feature vector for t
12: $l_t \leftarrow \lfloor \frac{K*(\tau_t - \tau_{min})}{(\tau_{max} - \tau_{min})} \rfloor$ ▷ Find which partition τ_t is in.
13: Output $< l_t, v_t >$

The algorithm first determines the ground-truth ranking of IR systems (Line 1). It then starts the process of generating the data records for each possible topic subset size for the targeted test collection (Line 2). For each topic subset size i, we randomly select W topic sets (each of size i, assuming that these are already-selected topics (Line 4)). For each topic t of the non-selected topics, we rank the systems over the randomly selected topics plus t (Lines 6-7). We next calculate Kendall's τ score achieved by the addition of each topic (Line 8). Subsequently, we divide the range between maximum and minimum τ scores into K equal ranges, each will represent a grade label. For example, let K be 10, maximum τ score be 0.9, and minimum one be 0.7, then, τ ranges for labeling will be [0.7-0.72), [0.72-0.74), ..., [0.88-0.9]. The labels for the topics are assigned from 0 to $(K - 1)$ based on which range their τ score fits in. For example, if we achieve 0.73 τ score for a particular topic, then the label for the corresponding data record will be 1. Finally, we compute feature vector of each topic, assign the labels, and output the data records for the current scenario (Lines 10-13). We repeat this process for W times (Line 3) to capture more different scenarios for the given topic subset size. We can also repeat the process using other test collections and merge the resulted data.

5 Evaluation

In this section, we evaluate our proposed L2R topic selection approach with respect to our research questions and baseline methods. Section 5.1 details our experimental setup, including generation of training data and tuning of our L2R model parameters. We present results of our topic selection experiments (RQ-1) in Section 5.2 and discuss the evaluation of the parameters of our approach in Section 5.3. In Section 5.4, we report the results of our experiments for intelligent topic selection with a fixed budget (RQ-2) and considering different parameters in the debate of NaD vs. WaS judging: varying judging speed (RQ-3), topic generation time (RQ-4), and judging error (RQ-5).

5.1 Setup

We adopt the MART [20] implementation of RankLib library[1] for building our L2R model. To tune MART parameters, we partition our data into disjoint training, tuning, and testing sets. We assume that the ground-truth ranking of systems is given by MAP@100.

Test Collections. We consider two primary criteria in selecting test collections to use in our experiments: (1) the collection should contain many topics, providing a fertile testbed for topic selection experimentation, and (2) the set of topics used in training, tuning, and testing should be disjoint to avoid over-fitting. To satisfy these criteria, we adopt the TREC-9 [25] and TREC-2001 [26] Web Track collections, as well as TREC-2003 [53] and TREC-2004 [52] Robust Track collections. Details of these test collections are presented in Table 3. Note that all four collections target ad-hoc retrieval. We use TREC-9 and TREC-2001 test collections to generate our training data.

Table 3: Test Collections Used in Experiments

Collection	Topics	Dataset	Runs	Judgments	# Relevant
TREC-9 Web Track	451-500	WT10g	104	70070	2617
TREC-2001 Web Track	501-550	WT10g	97	70400	3363
Robust2003	601-650, 50 difficult topics from 301-450	TREC disks 4&5 minus Congressional Records	78	128796	6074
Robust2004	301-450, 601-700	TREC disks 4&5 minus Congressional Records	110	311410	17412

Robust2003 and Robust2004 collections are particularly well-suited to topic selection experimentation since they have relatively more topics (100 and 249, respectively) than many other TREC collections. However, because topics of Robust2003 were repeated in Robust2004, we define a new test collection subset which excludes all Robust2003 topics from Robust2004, referring to this subset as Robust2004$_{149}$. We use Robust2003 and Robust2004$_{149}$ collections for tuning and testing. When testing on Robust 2003, we tune parameters on Robust2004$_{149}$, unless otherwise noted. Similarly, when testing on Robust2004$_{149}$, we tune parameters on Robust2003.

Generation of Training Data. We generate 100K data records for each topic set size from 0-49 (*i.e.*, $N = 50$ and $W = 100K$ in Algorithm 3) for TREC-9 and TREC-2001 and remove the duplicates. The label range is set to 0-49 (*i.e.*, $K = 50$ in Algorithm 3) since each of TREC-9 and TREC-2001 has 50 topics. We merge the data records generated from each test collection to form our final training data. We use this training data in our experiments unless otherwise stated.

Parameter Tuning. To tune MART parameters, we fixed the number of trees to 50 and varied the number of leaves from 2-50 with a step-size of 2. For each of those 25 considered configurations, we built a L2R model and selected 50 topics (the standard number of topics in TREC collections) using the tuning set. At each iteration of the topic selection process, we ranked the systems based on the topics selected thus far and calculated Kendall's τ rank

[1] https://sourceforge.net/p/lemur/wiki/RankLib/

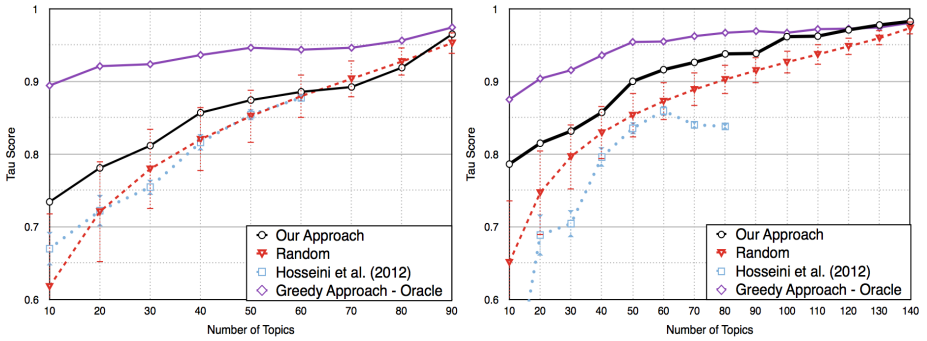

(a) Robust2003 (b) Robust2004$_{149}$

Figure 1: Selecting a fixed number of topics, using full pool judgments per topic, and evaluating with MAP.

correlation vs. the ground-truth system ranking. Finally, we selected the parameter configuration which achieved the highest average τ score while selecting the first 50 topics.

Baselines. We compare our approach to two baselines and an upper-bound oracle:

- **Baseline 1: Random**. For a given topic subset size M, randomly select topics R times and calculate the average Kendall's τ score achieved over the R trials.
- **Baseline 2: Hosseini et al. [27]**. We implemented Hosseini et al. [27]'s method (since no implementation is available from the authors). They do not specify SVM parameters used, so we adopt typical settings.
- **Upper-Bound: Greedy Oracle**. Algorithm 1 (defined in Section 4.1) selects the topic at each iteration that maximizes Kendall's τ rank correlation with the ground-truth ranking, assuming advance knowledge of relevance judgments for all topics. This serves as a useful oracle upper-bound, since in practice we would only collect judgments for a topic after it was selected.

5.2 Selecting A Fixed Number of Topics

In our first set of experiments, we evaluate our proposed L2R topic selection approach vs. baselines in terms of Kendall's τ rank correlation achieved as a function of number of topics (**RQ-1**). We assume the full pool of judgments are collected for each selected topic and used to evaluate systems. Due to the stochastic nature of the baselines, we run the random topic selection approach 10K times and Hosseini et al. [27]'s (more computationally-intensive) method 50 times, plotting the average performance of each method and its standard deviation.

Figure 1 shows results on Robust2003 and Robust2004$_{149}$ collections. Given the computational complexity of Hosseini et al.[27]'s method, which re-trains the classifier at each iteration, we could only select 63 topics for Robust2003 and 77 topics for Robust2004$_{149}$ after 2 days of execution[2], so its plots terminate early. The upper-bound *Greedy Oracle* is seen to achieve 0.90 τ score (a traditionally-accepted threshold for acceptable correlation [56]) with only 12 topics in Robust2003 and 20 topics in Robust2004$_{149}$. Our proposed L2R method strictly outperforms

[2]Two days is the time limit for executing programs on the computing cluster we used for experiments.

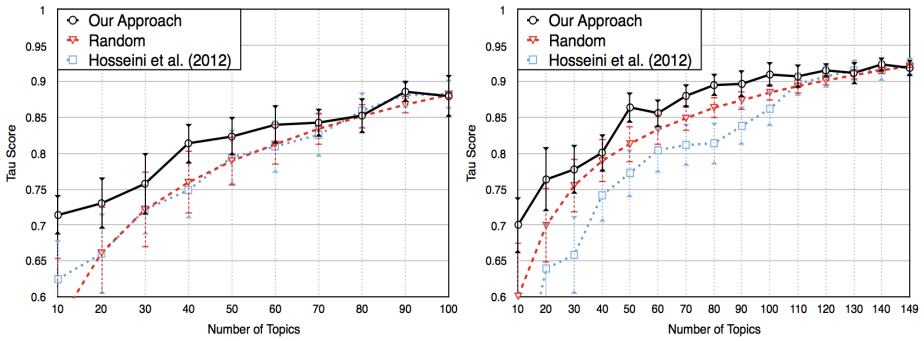

(a) Robust2003, 64 Judgments Per Topic

(b) Robust2004$_{149}$, 64 Judgments Per Topic

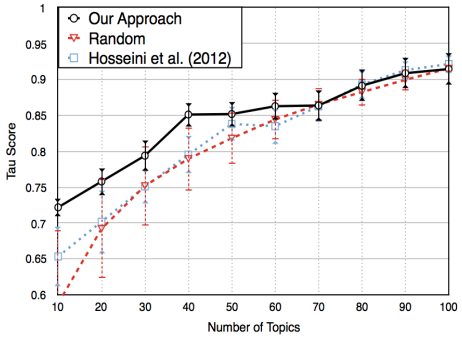

(c) Robust2003, 128 Judgments Per Topic

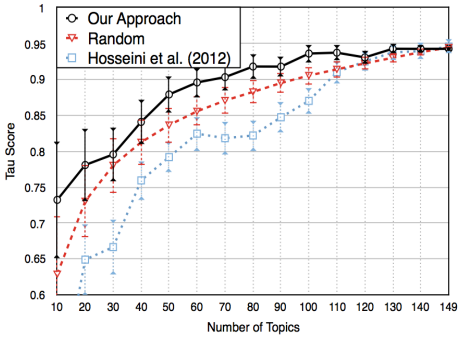

(d) Robust2004$_{149}$, 128 Judgments Per Topic

Figure 2: Selecting a fixed number of topics and a limited number of documents to judge per topic (using StatAP).

baselines for Robust2004$_{149}$ and outperforms baselines for Robust2003 except when 70% and 80% of topics are selected. Relative improvement over baselines is seen to increase as the number of topics is reduced. This suggests that our L2R method becomes more effective as either fewer topics are used, or as more topics are available to choose between when selecting a fixed number of topics.

In our next experiment, instead of assuming the full document pool is judged for each selected topic, we consider a more parsimonious judging condition in which statAP [39] is used to select only 64 or 128 documents to be judged for each selected topic. Because statAP is stochastic, we sample the documents to be judged 20 times for each selected topic and report average results.

The average τ scores for each method are shown in **Figure 2**. The vertical bars represent the standard deviation across trials. Overall, similar to the first set of experiments, our approach outperforms the baselines in almost all cases and becomes more preferable as the number of selected topics decreases. Similar to the previous experiment with full pooling, our L2R approach performs relatively weakest on Robust 2003 when 70 or 80 topics are selected for Robust2003. In this case, our L2R approach is comparable to random selection (with slight increase over it), whereas with the last experiment we performed slightly worse than random for 70 or 80 topics on this collection.

We were surprised to see Hosseini et al. [27]'s topic selection method performing worse than random in our experiments, contrary to their reported results. Consequently, we investigated this in great detail. In comparing results

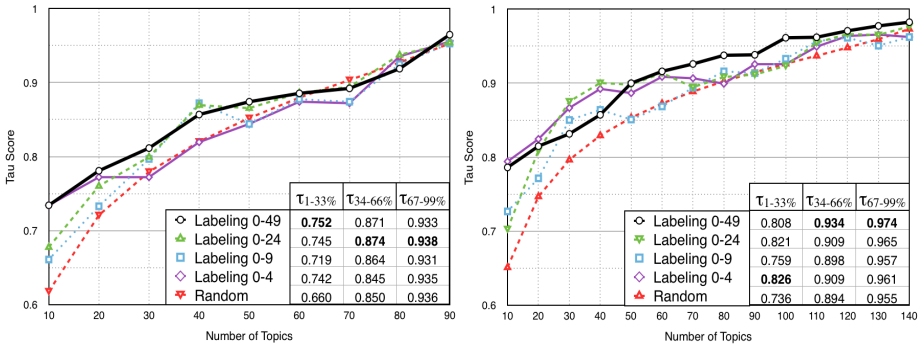

(a) Robust2003 (b) Robust2004$_{149}$

Figure 3: Effect of label ranges in training data for our learning-to-rank model.

.

of our respective random baselines, we noted that our own random baseline performed $\tau \approx 0.12$ better on average than theirs over the 20 results they report (using 10, 20, 30, ..., 200 topics), despite our carefully following their reported procedure for implementing the baseline. To further investigate this discrepancy in baseline performance, we also ran our random baseline on TREC-8 and compared our results with those reported by Guiver, Mizzaro, and Robertson [22]. Results were quite similar. Hosseini et al. [27] kindly discussed the issue with us, and the best explanation we could find was that they took, "special care when considering runs from the same participant", so perhaps different preprocessing of participant runs between our two studies may contribute to this empirical discrepancy.

5.3 Robustness and Parameter Sensitivity

The next set of experiments we report assess our L2R method's effectiveness across different training datasets and parameterizations. In addition to presenting results for all topics, we also compute the average τ score over 3 equal-sized partitions of the topics. For example, in Robust2004, we calculate the average τ scores for each of the following partitions: 1-50 (denoted by $\tau_{1-33\%}$), 51-100 (denoted by $\tau_{34-66\%}$) and 101-149 (denoted by $\tau_{67-100\%}$). These results are presented in a table within each figure.

Effect of Label Range in Training Set: As explained in Section 4.3.2, we can assign labels to data records in various ranges. In this experiment, we vary the label range parameter (K in Line 12 of Algorithm 3) and compare the performance of our approach with the corresponding training data on Robust2003 and Robust2004$_{149}$ test collections. The results are shown in **Figure 3**. It is hard to draw a clear conclusion since each labeling range has varying performances in different cases. For instance, when we use 5 labels only (*i.e.*, Labeling 0-4), it has very good performance with few selected topics. As the number of topics increases, its performance becomes very close to the random approach. Considering the results in Robust2003 and Robust2004$_{149}$ together, using 50 labels (*i.e.*, labeling 0-49) gives more consistent and better results than others. Using 25 labels are better than using 10 or 5 labels, in general. Therefore, we observe that *using fine grained labels yields better results with our L2R approach.*

Effect of Size of *Tuning* Dataset: In this experiment, we evaluate how robust our approach is to having fewer topics available for tuning. For this experiment, we randomly select 50 and 75 topics from Robust2003 and remove the not-selected ones from the test collection. We refer to these reduced tuning sets as R3(50) and R3(75). We

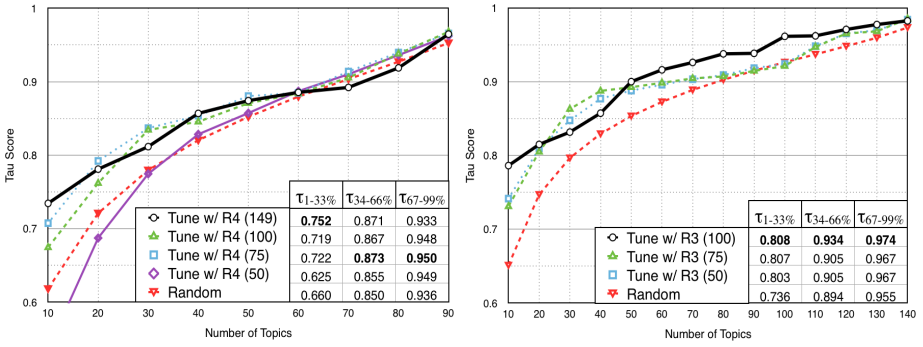

(a) Robust2003 (b) Robust2004$_{149}$

Figure 4: Effect of test collections used for parameter tuning.

.

use these reduced tuning sets for testing on Robust2004$_{149}$. When we test with Robust2003, we perform a similar approach. That is, we randomly select 50, 75, and 100 topics from Robust2004$_{149}$ and follow the same procedure. We repeat this process 5 times and calculate the average τ score achieved.

The results are presented in **Figure 4**. The vertical bars represent the standard deviation across 5 different trials. As expected, over Robust2004$_{149}$, we achieve the best performance when we tune with all 100 topics (*i.e.*, actual Robust2003); employing 75 topics is slightly better than employing 50 topics. Over Robust2003, when the number of selected topics is $\leq$33% of the whole topic pool size, tuning with 149 topics gives the best results. For the rest of the cases, tuning with 75 topics gives slightly better results than others. As expected, tuning with only 50 topics yields the worst results in general. Intuitively, *using test collections with more tuning topics is seen to yield better results.*

Effect of Test Collections Used in *Training*: In this experiment, we fix the training data set size, but vary the test collections used for generating the training data. For the experiments so far, we had generated 100K data records for each topic set size from 0-49 with TREC-9 and TREC-2001 and subsequently combined both (yielding 200K records in total). In this experiment, in addition to this training data, we generate 200K data records for each topic set size from 0-49 with TREC-9 and TREC-2001, and use them separately. That is, we have 3 different datasets (namely, T9&T1, T9 and T1) and each dataset has roughly the same number of data records. The results are shown in **Figure 5**. As expected, using more test collections leads to better and more consistent results. Therefore, instead of generating more data records from the same test collection, *diversifying the test collections in present in the training data appears to increase our L2R method's effectiveness*.

5.4 Topic Selection with a Fixed Budget

Next, we seek to compare NaD vs. WaS judging when topics are selected intelligently (**RQ-2**), considering also familiarization of assessors to topics (**RQ-3**), the effect of topic generation cost (**RQ-4**) and judging error (**RQ-5**). We evaluate the performance of the methods using statAP [39]. Documents are sampled 20 times for each case and average performance is reported. The budget is distributed equally among topics. In each experiment, we exhaust the full budget for the selected topics, *i.e.*, as the number of topics increases, the number of judgments per topic decreases, and vice-versa.

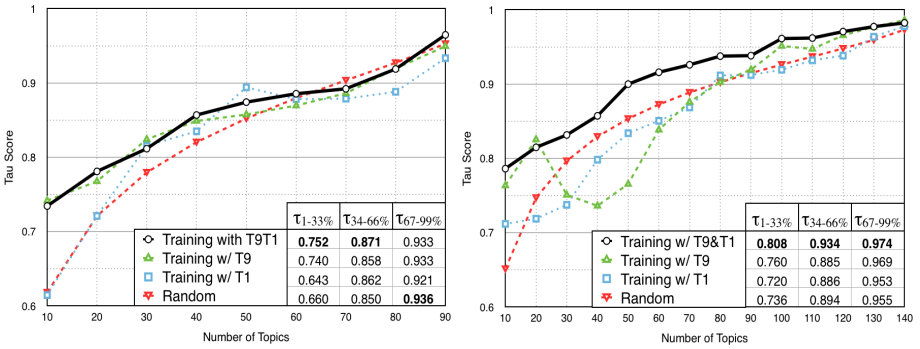

(a) Robust2003 (b) Robust 2004_{149}

Figure 5: Effect of test collections used for generating training data.

Effect of Familiarization to Topics when Judging: As discussed in Section 2.3, Carterette et al. [15] found that as the number of judgments per topic increases (when collecting 8, 16, 32, 64 or 128 judgments per topic), the median time to judge each document decreases respectively: 15, 13, 15, 11 and 9 seconds. Because prior work comparing NaD vs. WaS judging did not consider variable judging speed as a function of topic depth, we revisit this question, considering how faster judging with greater judging depth per topic may impact the tradeoff between deep vs. shallow judging in maximizing evaluation reliability for a given assessment time budget.

Using Carterette et al. [15]'s data points, we fit a piece-wise judging speed function (**Equation 9**) to simulate judging speed as a function of judging depth as illustrated in **Figure 6**. According to this model, judging a single document takes 15 seconds if there are 32 or fewer judgments per topic (*i.e.*, as the assessor "warms up"). After 32 judgments, the assessors become familiar with the topic and start judging faster. Because judging speed cannot increase forever, we assume that after 128 judgments, judging speed becomes stable at 9 seconds per judgment.

$$f(x) = \begin{cases} 15, & \text{if } x \leq 32 \\ 8.761 + 16.856 \times e^{-0.0316 \times x}, & \text{if } 32 < x < 127 \\ 9, & \text{otherwise} \end{cases} \quad (9)$$

For the constant judging case, we set the judging speed to 15 seconds per document. For instance, if our total budget is 100 hours and we have 100 topics, then we spend 1 hour per topic. If judging speed is constant, we judge $60\ min \times 60\frac{sec}{min} \div 15\frac{seconds}{judgment} = 240\ judgments$ for each topic. However, if judging speed increases according to our model, we can judge a larger set of 400 documents per topic in the same one hour.

We set our budget to 40 hours for both test collections and assume that generating the topics has no cost. We intelligently select topics using our proposed approach. Results are shown in **Figure 7**. We can see that the additional judgments due to faster judging results in a higher τ score after 30 topics, and its effect increases as the number of topics increases. Since the results are significantly different (p value of paired t-test is 0.0075 and 0.0001 in experiments with Robust2003 and Robust2004_{149}, respectively), results suggest that *topic familiarity (i.e., judging speed) impacts optimization of evaluation budget for deep vs. shallow judging*.

Effect of Topic Generation Cost: As discussed in Section 2, topic generation cost (TGC) can vary significantly. In order to understand TGC's effect, we perform intelligent topic selection with our L2R method and vary TGC from 76 seconds (*i.e.*, time needed to convert a query to topic, as reported in [15]) to 2432 seconds (*i.e.*, 32×76) in

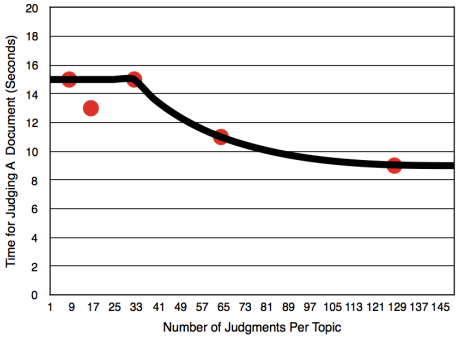

Figure 6: Illustration of Equation 9 modeling the judging time per document as a function of judgments per topic, as fit to actual data points reported by [15].

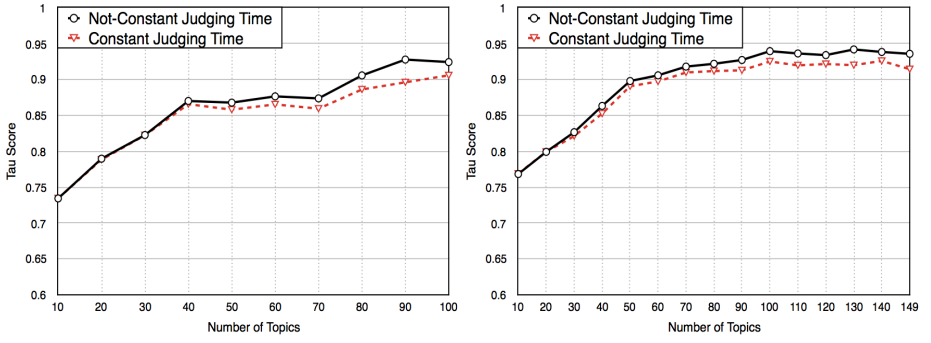

(a) Robust2003 (b) Robust 2004_{149}

Figure 7: Constant vs. Non-Constant Judging Speed. Topic selection is performed with our method. For each case, we run 20 times and report average τ score.

geometric order while fixing the budget to 40 hours for both test collections. Note that TREC spends 4 hours to generate a final topic [1], which is almost 5 times more than 2432 seconds. We assume that judging speed is constant (*i.e.*, 15 seconds per judgment). For instance, if TGC is 76 seconds and we select 50 topics, then we subtract $50 \times 76 = 3800$ seconds from the total budget and use the remaining for document judging.

The results are shown in **Figure 8**. When the topic generation cost is $\leq$152 seconds, results are fairly similar. However, when we spend more time on the topic generation, after selecting a number of topics, τ scores achieved starts decreasing due to insufficient budget left for judging the documents. In Robust2004_{149}, the total budget is not sufficient to generate more than 118 topics when TGC is 1216 seconds. Therefore, no judgment can be collected when the number of topics is 120 or higher. Considering the results for Robust2004_{149} test collection, when TGC is 304 seconds (which is close to that mentioned in [14]), we are able to achieve better performance with 80-110 topics instead of employing all topics. When TGC is 608 seconds, using only 50 topics achieves higher τ score than employing all topics (0.888 vs. 0.868). Overall, the results suggest that *as the topic generation cost increases, NaD judging becomes more cost-effective than WaS judging.*

Another effect of topic generation cost can be observed in the *reliability of the judgments*, as discussed in

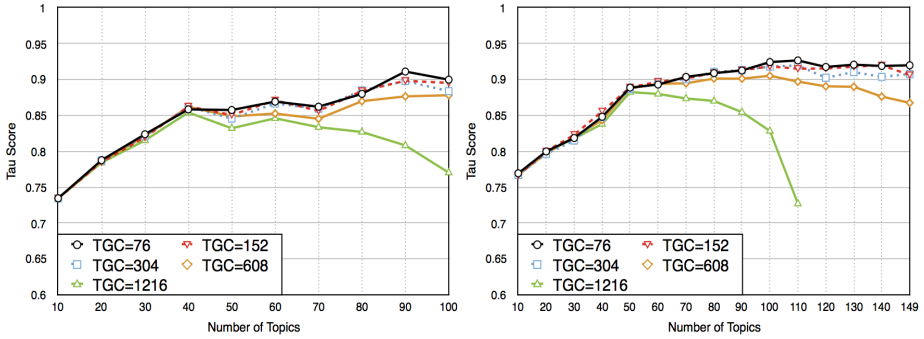

(a) Robust 2003 (b) Robust 2004_{149}

Figure 8: Effect of Topic Generation Cost (times listed in seconds). For all cases, we apply statAP and sample the documents 20 times. We assume that each judgment requires 15 seconds. We use our method to select the topics.

Section 2. In this experiment, we consider a scenario where the assessors rely on the topic definitions and poorly-defined topics can cause inconsistent judgments. In order to simulate this scenario, we assume that 8% of judgments are inconsistent when $TGC = 76$ seconds. The accuracy of judgments increases by 2% as TGC doubles. So when TGC = 1216 seconds, assessors can achieve perfect judgments. Note that the assessors in this scenario are much more reliable than what is reported in [33]. In order to implement this scenario, we randomly flip over judgments of *qrels* based on the corresponding accuracy of judging. The ground-truth ranking of the systems are based on the original judgments. We set the total budget to 40 hours and assume that judging a single document takes 15 seconds. We use our method to select the topics. We repeat the process 50 times and report the average. The results are shown in **Figure 9**. In Robust 2003, the achieved Kendall's τ scores increase as TGC increases from 76 to 608 seconds due to more consistent judging (opposite of what we observe in Figure 8). In Robust2004_{149}, when the number of topics is 100 or less, Kendall's τ score increases as TGC increases from 76 to 608. However, when the number of topics is more than 100, τ scores achieved with $TGC = 608$ case start decreasing due to insufficient amount of budget for judging. We observe a similar pattern in Robust2003 with $TGC = 1216$ case: We achieve highest τ scores with $TGC = 1216$ case when number of topics is 60 or less. However, its performance starts decreasing later on.

In general, the results suggest that poorly-defined topics should be avoided if they have negative effect on the consistency of the relevance judgments. However, spending more time to develop high quality topics can significantly increase the cost. Therefore, *NaD becomes preferable over WaS when we target constructing high quality topics*.

Varying budget: In this set of experiments, we vary our total budget from 20 to 40 hours. We assume that the assessors judge faster as they judge more documents, up to a point, based on our model given in Equation 9. We also assume that topic generation cost is 76 seconds. We run the random approach 1000 times and sample documents 20 times for each case. The results are shown in **Figure 10**. In Robust2004_{149}, our approach is better than random method in all cases. In Robust2003, our approach outperforms random method in the selection of the first 50 topics, while their performances become similar when we select more than 50 topics. When our budget is 20 hours, we are able to achieve higher τ scores by reducing the number of topics in both test collections. In Robust2004_{149}, when our budget is 30 hours, using 90-130 topics leads to higher τ scores than using all topics. When our budget is 40 hours, we are able to achieve similar τ scores by reducing the number of topics to 100 topics. Note that we are never

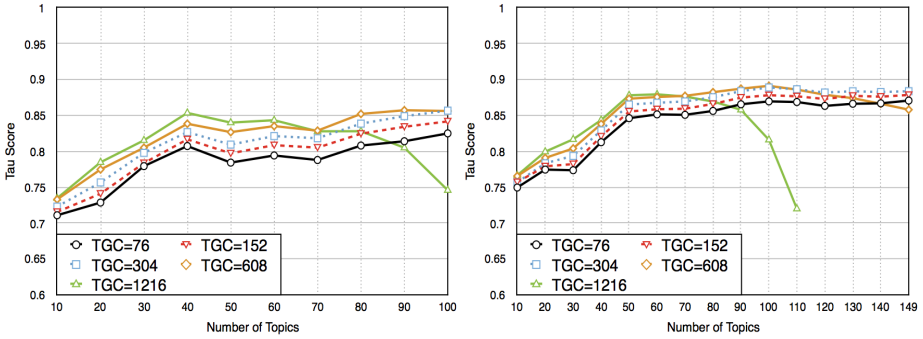

(a) Robust 2003 (b) Robust 2004 with 149 Topics

Figure 9: Effect of topic generation cost with imperfect judgments. For all cases, we apply statAP and sample the documents 20 times. We repeat the process 50 times and report the average overall. We assume that each judgment requires 15 seconds. We use our method to select the topics.

able to achieve higher τ scores by reducing the number of topics in the random method. That is to say that, *WaS judging leads to a better ranking of systems if we select the topics randomly, as reported by other studies [57, 48, 9]; however, if we select the topics more intelligently, we can achieve a better ranking by reducing the number of topics for a given budget.*

Re-usability of Test Collections: In this experiment, we compare our approach with the random topic selection method in terms of re-usability of the test collections with the selected topics. We again set topic generation cost to 76 seconds and assume non-constant judging speed. We vary the total budget from 20 hours to 40 hours, as in the previous experiment. In order to measure the re-usability of the test collections, we adopt the following process. For each topic selection method, we first select the topics for the given topic subset size. Then, by using only the selected topics, we apply the leave-one-group-out method [54]: for each group, we ignore the documents which only that group contributes to the pool and sample documents based on only remaining documents. Then, statAP score is calculated for the runs of the corresponding group. After applying this for all groups, we rank the systems based on their statAP scores and calculate Kendall's τ score compared to the ground-truth ranking of the retrieval systems. We repeat this process 20 times for our method and 5000 times for random method by re-selecting the topics. The results are shown in **Figure 11**. The vertical bars represent the standard deviation and the dashed horizontal line represents the performance when we employed all topics. There are several observations we can make from the results. First, our proposed method yields more *re-usable* test collections than random method in almost all cases. As the budget decreases, our approach becomes more effective in order to construct reusable test collections. Second, in all budget cases for both test collections, we can reach same/similar re-usability scores with fewer topics. Lastly, by using the random topic selection method, we can never construct a more reusable test collection by reducing the number of topics. However, by intelligently selecting the topics, we can increase the reusability of the test collections in all budget scenarios for Robust2004_{149} and 20-hours-budget scenario for Robust2003. Therefore, *the results suggest that NaD judging can yield more reusable test collections than WaS judging, when topics are selected intelligently.*

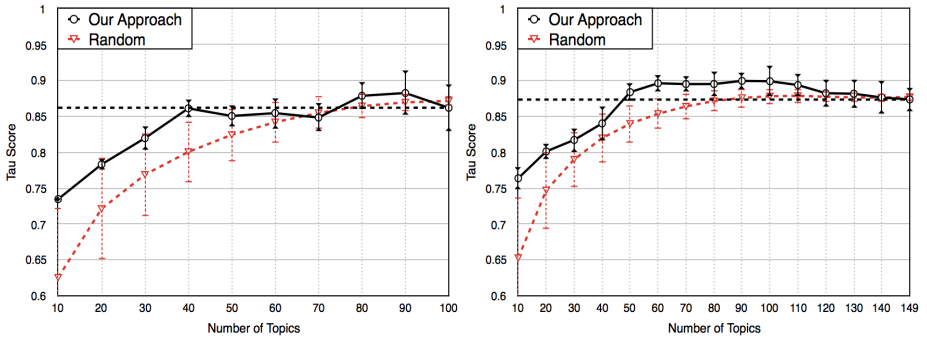

(a) Robust 2003 - 20 Hours

(b) Robust 2004_{149} - 20 Hours

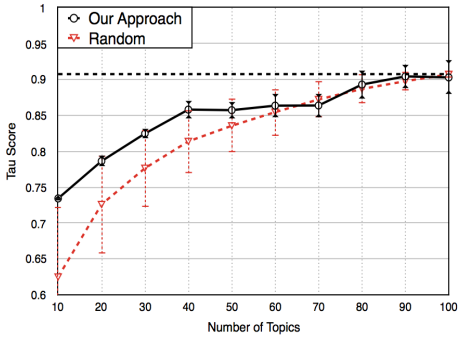

(c) Robust 2003 - 30 Hours

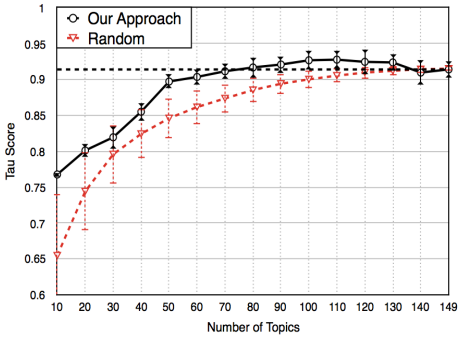

(d) Robust 2004_{149} - 30 Hours

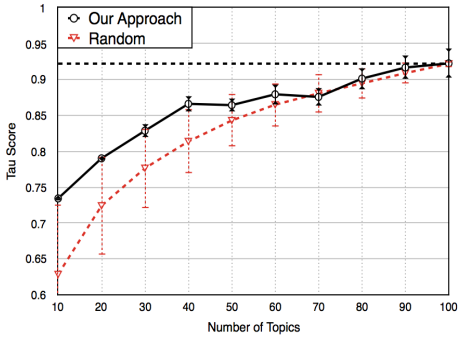

(e) Robust 2003 - 40 Hours

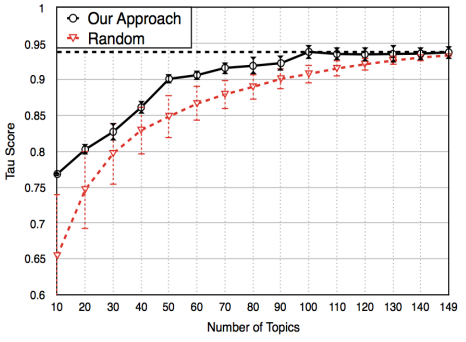

(f) Robust 2004_{149} - 40 Hours

Figure 10: Topic selection with different test collections when evaluation metric is statAP and we have a fixed budget. For random method, the number of query selection trials is 1000. For both method, we run statAP 20 times and took the average of all trials. We assume that each judgment requires 15 seconds. The dashed horizontal line represents the performance when we employed all topics.

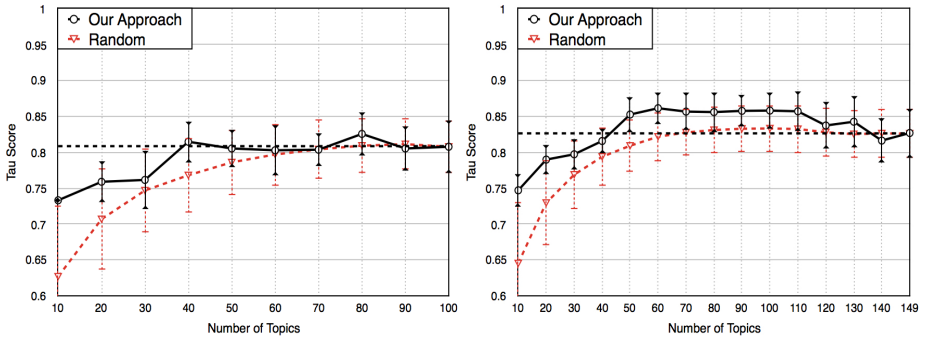

(a) Robust 2003 - 20 Hours

(b) Robust 2004_{149} - 20 Hours

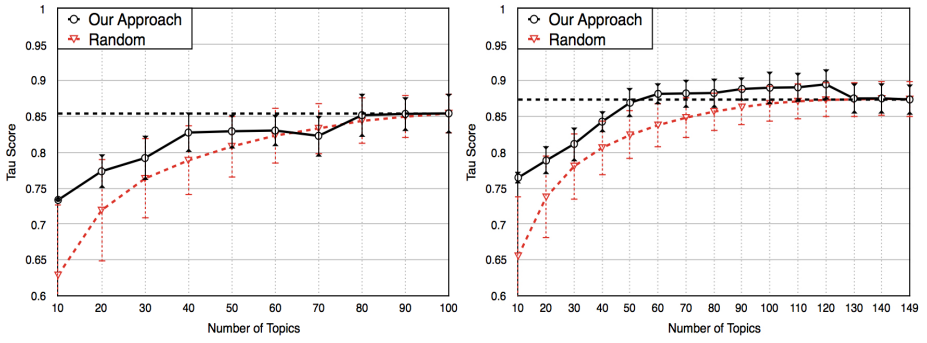

(c) Robust 2003 - 30 Hours

(d) Robust 2004_{149} - 30 Hours

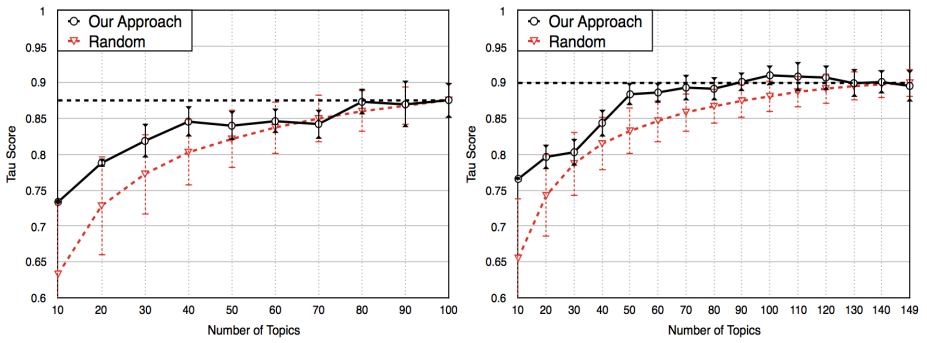

(e) Robust 2003 - 40 Hours

(f) Robust 2004_{149} - 40 Hours

Figure 11: Re-usability Performance. For our method, we run statAP 20 times and take the average of tau scores achieved. For random approach, we select queries 5000 times and calculate statAP once for each query set. Topic generation cost is set to 76 seconds and we assume that the judging speed increases as they judge more documents. The vertical bars represents the standard deviation. The dashed horizontal line represents the performance when we employed all topics.

6 Conclusion

In this paper, we proposed a learning-to-rank based topic selection method in order to find a *good* subset of topics for an IR test collection. We developed 63 features to represent interaction between IR systems and topics, and generated training data from existing test collections using a novel approach. We employed the L2R model to select topics in an iterative approach. We evaluated our approach in various cases and compared it against two state-of-the-art baselines. We also revisited deep vs. shallow judging debate, investigating many factors including intelligent selection of topics, increased judging speed when judging more documents for the same topic, and the effect of topic generation cost on the budget utilization and the quality of relevance judgments.

The main observations from our extensive evaluation are as follows. First, in almost all cases, our proposed approach selects better topics yielding more reliable evaluation than other baselines. Second, shallow judging is preferable than deep judging if topics are selected randomly, confirming findings of prior work. However, when topics are selected intelligently, deep judging often achieves greater evaluation reliability for the same evaluation budget than shallow judging. Third, assuming that judging speed increases as more documents for the same topic are judged, increased judging speed has significant effect on evaluation reliability, suggesting that it should be another parameter to be considered in deep vs. shallow judging trade-off. Fourth, as topic generation cost increases, deep judging becomes preferable than shallow judging. Finally, assuming that short topic generation times reduce the quality of topics, and thereby relevance judgments consistency, it is better to increase quality of topics and collect fewer judgments instead of collecting many judgments with low-quality topics. This also makes deep judging preferable than shallow judging in many cases, due to increased topic generation cost.

As future work, we plan to conduct qualitative analysis to understand why some set of topics are better than others in terms of evaluating retrieval systems. We believe that understanding this will be useful to intelligently design topic sets and increase evaluation reliability.

Acknowledgments

This work was made possible by NPRP grant# NPRP 7-1313-1-245 from the Qatar National Research Fund (a member of Qatar Foundation). The statements made herein are solely the responsibility of the authors. We thank the Texas Advanced Computing Center (TACC) at the University of Texas at Austin for computing resources enabling this research.

References

[1] Personal communication with Ellen Voorhees at NIST.

[2] James Allan, Javed A. Aslam, Ben Carterette, Virgil Pavlu, and Evangelos Kanoulas. "Million query track 2008 overview". In: In Proceedings of the Seventeenth Text REtrieval Conference (TREC 2007. 2008.

[3] James Allan, Ben Carterette, Javed A Aslam, Virgil Pavlu, Blagovest Dachev, and Evangelos Kanoulas. Million query track 2007 overview. Tech. rep. DTIC Document, 2007.

[4] Omar Alonso and Stefano Mizzaro. "Can we get rid of TREC assessors? Using Mechanical Turk for relevance assessment". In: Proceedings of the SIGIR 2009 Workshop on the Future of IR Evaluation. Vol. 15. 2009, p. 16.

[5] Javed A Aslam, Virgil Pavlu, and Emine Yilmaz. "A statistical method for system evaluation using incomplete judgments". In: Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval. ACM. 2006, pp. 541548.

[6] Javed A Aslam and Emine Yilmaz. "Inferring document relevance from incomplete information". In: Proceedings of the sixteenth ACM conference on Conference on information and knowledge management. ACM. 2007, pp. 633642.

[7] David Banks, Paul Over, and Nien-Fan Zhang. "Blind men and elephants: Six approaches to TREC data". In: Information Retrieval 1.1-2 (1999), pp. 734.

[8] Andrea Berto, Stefano Mizzaro, and Stephen Robertson. "On using fewer topics in information retrieval evaluations". In: Proceedings of the 2013 Conference on the Theory of Information Retrieval. ACM. 2013, p. 9.

[9] David Bodoff and Pu Li. "Test theory for assessing IR test collections". In: Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval. ACM. 2007, pp. 367374.

[10] Chris Buckley, Darrin Dimmick, Ian Soboroff, and Ellen Voorhees. "Bias and the limits of pooling". In: Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval. ACM. 2006, pp. 619620.

[11] Chris Buckley and Ellen M Voorhees. "Evaluating evaluation measure stability". In: Proceedings of the 23rd annual international ACM SIGIR conference on Research and development in information retrieval. ACM. 2000, pp. 3340.

[12] Chris Buckley and Ellen M Voorhees. "Retrieval evaluation with incomplete information". In: Proceedings of the 27th annual international ACM SIGIR conference on Research and development in information retrieval. ACM. 2004, pp. 2532.

[13] Ben Carterette, James Allan, and Ramesh Sitaraman. "Minimal test collections for retrieval evaluation". In: Proceedings of the 29th annual international ACM SIGIR conference on Research and development in information retrieval. ACM. 2006, pp. 268275. 27

[14] Ben Carterette, Virgil Pavlu, Evangelos Kanoulas, Javed A Aslam, and James Allan. "Evaluation over thousands of queries". In: Proceedings of the 31st annual international ACM SIGIR conference on Research and development in information retrieval. ACM. 2008, pp. 651658.

[15] Ben Carterette, Virgil Pavlu, Evangelos Kanoulas, Javed A Aslam, and James Allan. "If I had a million queries". In: European conference on information retrieval. Springer. 2009, pp. 288300.

[16] Ben Carterette and Mark D Smucker. "Hypothesis testing with incomplete relevance judgments". In: Proceedings of the sixteenth ACM conference on Conference on information and knowledge management. ACM. 2007, pp. 643652.

[17] Cyril W Cleverdon. "The evaluation of systems used in information retrieval". In: Proceedings of the international conference on scientific information. Vol. 1. National Academy of Sciences Washington, DC, 1959, pp. 687698.

[18] Gordon V Cormack, Christopher R Palmer, and Charles LA Clarke. "Efficient construction of large test collections". In: Proceedings of the 21st annual international ACM SIGIR conference on Research and development in information retrieval. ACM. 1998, pp. 282289.

[19] Ronan Cummins, Joemon Jose, and Colm O'Riordan. "Improved query performance prediction using standard deviation". In: Proceedings of the 34th international ACM SI- GIR conference on Research and development in Information Retrieval. ACM. 2011, pp. 10891090.

[20] Jerome H Friedman. "Greedy function approximation: a gradient boosting machine". In: Annals of statistics (2001), pp. 11891232.

[21] Catherine Grady and Matthew Lease. "Crowdsourcing document relevance assessment with mechanical turk". In: Proceedings of the NAACL HLT 2010 workshop on creating speech and language data with Amazons mechanical turk. Association for Computational Linguistics. 2010, pp. 172179.

[22] John Guiver, Stefano Mizzaro, and Stephen Robertson. "A few good topics: Experiments in topic set reduction for retrieval evaluation". In: ACM Transactions on Information Systems (TOIS) 27.4 (2009), p. 21.

[23] Claudia Hauff, Djoerd Hiemstra, Leif Azzopardi, and Franciska De Jong. "A case for automatic system evaluation". In: European Conference on Information Retrieval. Springer. 2010, pp. 153165.

[24] Claudia Hauff, Djoerd Hiemstra, Franciska De Jong, and Leif Azzopardi. "Relying on topic subsets for system ranking estimation". In: Proceedings of the 18th ACM conference on Information and knowledge management. ACM. 2009, pp. 18591862.

[25] David Hawking. "Overview of the TREC-9 Web Track." In: TREC. 2000.

[26] David Hawking and Nick Craswell. "Overview of the TREC-2001 web track". In: NIST special publication (2002), pp. 6167.

[27] Mehdi Hosseini, Ingemar J Cox, Natasa Milic-Frayling, Milad Shokouhi, and Emine Yilmaz. "An uncertainty-aware query selection model for evaluation of IR systems". In: Proceedings of the 35th international ACM SIGIR conference on Research and development in information retrieval. ACM. 2012, pp. 901910.

[28] Mehdi Hosseini, Ingemar J Cox, Natasa Milic-Frayling, Vishwa Vinay, and Trevor Sweeting. "Selecting a subset of queries for acquisition of further relevance judgements". In: Conference on the Theory of Information Retrieval. Springer. 2011, pp. 113124.

[29] K. Spa rck Jones and C. J. van Rijsbergen. "Report on the need for and provision of an "ideal" information retrieval test collection (British Library Research and Development Report No. 5266)". In: (1975), p. 43.

[30] Gabriella Kazai and Homer Sung. "Dissimilarity based query selection for efficient preference based IR evaluation". In: European Conference on Information Retrieval. Springer. 2014, pp. 172183.

[31] Maurice G Kendall. "A new measure of rank correlation". In: Biometrika 30.1/2 (1938), pp. 8193.

[32] Ting-Chu Lin and Pu-Jen Cheng. "Query sampling for learning data fusion". In: Proceedings of the 20th ACM international conference on Information and knowledge management. ACM. 2011, pp. 141146.

[33] Tyler McDonnell, Matthew Lease, Mucahid Kutlu, and Tamer Elsayed. "Why Is That Relevant? Collecting Annotator Rationales for Relevance Judgments". In: Proceedings of the 4th AAAI Conference on Human Computation and Crowdsourcing (HCOMP). AAAI. 2016, pp. 139148.

[34] Rishabh Mehrotra and Emine Yilmaz. "Representative & Informative Query Selection for Learning to Rank using Submodular Functions". In: Proceedings of the 38th International ACM SIGIR Conference on Research and Development in Information Retrieval. ACM. 2015, pp. 545554.

[35] Stefano Mizzaro and Stephen Robertson. "Hits hits TREC: exploring IR evaluation results with network analysis". In: Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval. ACM. 2007, pp. 479 486.

[36] Alistair Moffat, William Webber, and Justin Zobel. "Strategic system comparisons via targeted relevance judgments". In: Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval. ACM. 2007, pp. 375382.

[37] Shiva Imani Moghadasi, Sri Devi Ravana, and Sudharshan N Raman. "Low-cost evaluation techniques for information retrieval systems: A review". In: Journal of Informetrics 7.2 (2013), pp. 301312.

[38] Rabia Nuray and Fazli Can. "Automatic ranking of information retrieval systems using data fusion". In: Information Processing & Management 42.3 (2006), pp. 595614.

[39] V Pavlu and J Aslam. A practical sampling strategy for efficient retrieval evaluation. Tech. rep. Technical Report, College of Computer and Information Science, Northeastern University, 2007.

[40] Stephen Robertson. "On the contributions of topics to system evaluation". In: European conference on information retrieval. Springer. 2011, pp. 129140.

[41] Tetsuya Sakai. "Alternatives to bpref". In: Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval. ACM. 2007, pp. 7178.

[42] Tetsuya Sakai. "Designing test collections for comparing many systems". In: Proceedings of the 23rd ACM International Conference on Conference on Information and Knowledge Management. ACM. 2014, pp. 6170.

[43] Tetsuya Sakai. "Designing test collections that provide tight confidence intervals". In: Forum on information technology 2014 RD. Vol. 3. 2014.

[44] Tetsuya Sakai. "Topic set size design". In: Information Retrieval Journal (2015), pp. 1 28.

[45] Tetsuya Sakai. "Topic Set Size Design with Variance Estimates from Two-Way ANOVA." In: EVIA@ NTCIR. Citeseer. 2014.

[46] Mark Sanderson. Test collection based evaluation of information retrieval systems. Now Publishers Inc, 2010.

[47] Mark Sanderson, Falk Scholer, and Andrew Turpin. "Relatively relevant: Assessor shift in document judgements". In: Proceedings of the Australasian Document Computing Symposium. 2010, pp. 6067.

[48] Mark Sanderson and Justin Zobel. ”Information retrieval system evaluation: effort, sensitivity, and reliability”. In: Proceedings of the 28th annual international ACM SIGIR conference on Research and development in information retrieval. ACM. 2005, pp. 162 169.

[49] Falk Scholer, Diane Kelly, Wan-Ching Wu, Hanseul S Lee, and William Webber. ”The effect of threshold priming and need for cognition on relevance calibration and assessment”. In: Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval. ACM. 2013, pp. 623632.

[50] Ian Soboroff, Charles Nicholas, and Patrick Cahan. ”Ranking retrieval systems without relevance judgments”. In: Proceedings of the 24th annual international ACM SIGIR conference on Research and development in information retrieval. ACM. 2001, pp. 6673.

[51] Julián Urbano, Mónica Marrero,and Diego Martín.”On the measurement of test collection reliability”. In: Proceedings of the 36th international ACM SIGIR conference on Research and development in information retrieval. ACM. 2013, pp. 393402.

[52] Ellen M Voorhees. ”Overview of the TREC 2004 Robust Track.” In: TREC. Vol. 4. 2004.

[53] Ellen M Voorhees. ”Overview of TREC 2003.” In: TREC. 2003, pp. 113.

[54] Ellen M Voorhees. ”The philosophy of information retrieval evaluation”. In: Workshop of the Cross-Language Evaluation Forum for European Languages. Springer. 2001, pp. 355 370.

[55] Ellen M Voorhees. ”Topic set size redux”. In: Proceedings of the 32nd international ACM SIGIR conference on Research and development in information retrieval. ACM. 2009, pp. 806807.

[56] Ellen M Voorhees. ”Variations in relevance judgments and the measurement of retrieval effectiveness”. In: Information processing & management 36.5 (2000), pp. 697716.

[57] William Webber, Alistair Moffat, and Justin Zobel. ”Statistical power in retrieval experimentation”. In: Proceedings of the 17th ACM conference on Information and knowledge management. ACM. 2008, pp. 571580.

[58] Emine Yilmaz and Javed A Aslam. ”Estimating average precision when judgments are incomplete”. In: Knowledge and Information Systems 16.2 (2008), pp. 173211.

[59] Emine Yilmaz and Javed A Aslam. ”Estimating average precision with incomplete and imperfect judgments”. In: Proceedings of the 15th ACM international conference on In- formation and knowledge management. ACM. 2006, pp. 102111.

[60] Justin Zobel. ”How reliable are the results of large-scale information retrieval experiments?” In: Proceedings of the 21st annual international ACM SIGIR conference on Re- search and development in information retrieval. ACM. 1998, pp. 307314.